Binaryworks
Formerly DrupalPartners

AN EBOOK FOR DRUPAL ORGANIZATIONS

Turning Google's May 2026 AI Search Guidance into Visibility for Drupal Organizations

Cited, Not Displaced

How to make AI-driven search your highest-intent traffic channel, using the platform you already run.

CONTENTS

Inside This eBook

A practical, sequenced guide: from the problem, to the fix, to measurement, to a 90-day plan you can run on your existing Drupal platform.

— *Introduction: The Quiet Rewrite of Search*

01 **The New Reality of Search**

02 **What Google Actually Said**

03 **The One Shift: From Publishable to Citable**

04 **Fix Area One: Be Retrievable**

05 **Fix Area Two: Be Citable**

06 **Fix Area Three: Earn the Visit**

07 **What Google Always Rewards**

08 **Stop Funding the Noise**

09 **Measure What Changed**

10 **Your 90-Day Roadmap**

— *Your Next Step*

— *About BinaryWorks*

INTRODUCTION

The Quiet Rewrite of Search

For most of the last two decades, the goal of search visibility was simple to state: rank on page one. A high ranking meant clicks, and clicks meant traffic. That contract has quietly been rewritten.

Today, more than two billion people each month encounter an AI-generated answer at the top of Google Search before they ever see a list of links. Conversational AI search, where a person asks a question and receives a synthesized response, has moved from novelty to habit. The answer, not the link, has become the destination.

For the organizations that depend on being found (universities, healthcare systems, nonprofits, government agencies, publishers, and B2B enterprises), this raises an uncomfortable question. When an AI system answers on your behalf, is your organization the source it names, or the source it quietly skips?

This ebook is built to answer that question and to fix it. It draws on Google's first official guidance on AI search, published in May 2026, and translates it into a concrete plan for organizations running Drupal. The most reassuring finding is worth stating up front: almost everything you need is already within reach of your current platform. You are not facing a rebuild. You are facing a reconfiguration.

KEY INSIGHT

AI search did not introduce a separate algorithm to outsmart. It raised the bar on the fundamentals: being found, being clear, and being trustworthy. The organizations that win treat AI visibility as a discipline of structure, not a race for volume.

Prefer to watch first?

This ebook builds on our on-demand session for Drupal teams, which walks through every fix with live examples.

Watch the Webinar Recording

CHAPTER 01

The New Reality of Search

Search changed along four fronts at once. Understanding them is the foundation for everything that follows.

AI Overviews

Google now synthesizes an answer at the top of the results page, assembled from sources it has indexed and trusts. Most users read it and stop there.

AI Mode

A fully conversational experience. People ask, the AI answers, and no site visit is required to satisfy the question.

Zero-click search

More than 60 percent of searches now end without a single click. When the answer is served on the results page, the page becomes the endpoint rather than a waypoint.

Conversational queries

People no longer type clipped keywords. They ask full, specific questions, which means intent is richer, and the content that answers precisely is the content that surfaces.

The practical consequence is a fork. Every query your audience asks now resolves into one of two outcomes.

	Displaced	Cited
What happens	The AI answers; your link goes unclicked	You are the named source, and the visit arrives
Your position	Absent from the decision before a visit begins	Present, credited, and trusted
Traffic quality	Lost entirely	Higher intent; the reader arrives informed
Root cause	Pages Google cannot find, read, or trust	Visible, organized, first-hand content

> KEY INSIGHT
>
> The gap between displaced and cited is structural, not creative. It is rarely a matter of writing better. It is a matter of being retrievable, extractable, and attributable. That is good news, because structure is fixable.

CHAPTER 02

What Google Actually Said

In May 2026, after months of organizations rewriting pages "for AI," building new file types to feed AI systems, and hiring specialists for answer-engine and generative-engine optimization, Google published its first official guidance. The verdict landed hard: most of that effort was unnecessary.

The central statement was disarmingly plain. Optimizing for AI search is still SEO. There is no separate algorithm and no new system to game.

Consider what happens when a prospective student searches "best universities for environmental policy research." Google retrieves pages it has already indexed and ranked for that topic. The question then fans out into related searches such as faculty expertise, program outcomes, and research funding, and the AI assembles its answer from the pages it can find for each. Any page Google cannot find or rank is never a candidate to be cited. The same sequence runs whether the searcher is a student, a donor, a citizen, or a patient.

Google grouped its direction into three things to do and one thing to stop.

Area	What Google Asked For
Technical	Every page must be indexed, crawlable, and eligible to appear with a snippet
Content	Create unique, first-hand, non-commodity content; organize it clearly; add images and video
Experience	Deliver good page experience across devices, with clear, readable structure
Stop	A short list of specific tactics named as unnecessary for AI search

> THE CHAIN THAT MATTERS
>
> Not retrievable, therefore not citable, therefore not visible. If a page cannot clear the first gate, nothing downstream can save it.

CHAPTER 03

The One Shift: From Publishable to Citable

If you remember one idea from this ebook, remember this one. The new standard is not whether content is **publishable**. It is whether content is **citable**: something an AI system can extract, trust, and attribute to you.

Most organizations have plenty of publishable content. Far fewer have citable content. The difference is precise.

Publishable	Citable
Exists on a live page	Indexed and accessible to Googlebot
Written by someone at your organization	Carries a named author with stated expertise
Covers the topic	Answers a specific question in the first sentence
Has good formatting	Uses labeled sections an AI can extract cleanly
Published on time	Updated, with a visible date

A 30-second test you can run right now

Open any important page on your site and apply four checks:

1. Strip away the navigation and the page title. Can an AI still tell what this page is and who it is for?
2. Is the core answer in the first sentence, or buried in the fourth paragraph?
3. Is your organization named as the authority, or is the page effectively anonymous?
4. Would the meaning survive without the surrounding site design?

> RED FLAG
>
> If the answer is "no" to any of these, that page is, for practical purposes, invisible to AI search, no matter how good it looks to a human visitor.

CHAPTER 04

Fix Area One: Be Retrievable

Retrievability comes first because it is the precondition for everything else. A page that cannot be found cannot be read, and a page that cannot be read cannot be cited.

Four technical issues account for the majority of avoidable losses.

The indexability gate

A visibility block left over from a development environment can silently hide an entire section of a site. Consider a government agency that publishes new benefit and eligibility pages, then sees zero Search Console impressions six weeks later, because a "no-index" rule carried over from staging was never removed. The pages exist, but to search they do not. Every page that matters must be indexed before it can be cited.

Crawl budget

When a site carries forty thousand pages of old event listings, archives, and parameter URLs, Googlebot spends most of its time on pages that will never be cited, and may never reach the ones that should. Spend that finite crawl budget on the pages that earn citations.

What Google sees versus what you see

A page can look complete to a visitor while presenting a blank container to Googlebot, because the meaningful content is assembled by JavaScript after the fact. For content that must be found, server-rendered HTML is far more reliable.

One page, one address

When the same content lives at several URLs, because a page was rebuilt and old versions were never redirected, authority splits and no single version accumulates enough signal to win. Consolidate to one canonical address.

Issue	The Risk Chain	The Core Fix
Indexability	Page exists, not indexed, not retrieved, not cited	Remove dev-era blocks; confirm every key page is indexable; submit to Google
Crawl budget	Low-value pages crawled, budget wasted, key pages missed	De-index archives and utility pages; consolidate; submit a lean sitemap
JS rendering	Loads for visitors, blank for Googlebot, not indexed	Move critical content to server-rendered HTML; test with JavaScript disabled
Canonicals	Same content, many URLs, authority split, no page wins	One permanent canonical URL per topic; redirect duplicates

QUICK WIN

Take your five most important pages and view them with JavaScript disabled. If the substance disappears, Googlebot may not be seeing it either. That single check often explains a citation problem no amount of writing would have fixed.

CHAPTER 05

Fix Area Two: Be Citable

This is where most organizations have the largest hidden advantage, and the largest unrealized loss. The expertise usually already exists. It is simply locked inside formats and structures that AI cannot extract.

Your expertise is already there

Picture a nonprofit with fifteen years of program research and hundreds of documented outcomes, all sitting in PDFs, newsletters, and unstructured pages. A peer with comparable results, but structured and crawlable, gets cited in every AI answer. The expertise gap between them is zero. The structure gap is everything.

Lead with first-hand content

AI rewards what only your organization can say. Generic explanation is everywhere and easily replaced. Original, attributable detail (your own outcomes, your own data, your own named experts) is rare and valuable. Begin each priority page with a direct answer, name the source, and add real evidence.

Organize for extraction

A page can be accurate, detailed, and current, yet still fail, because every answer is buried inside paragraphs of running prose. Rewrite key pages with labeled sections, move eligibility, process steps, and outcomes into discrete fields, and lead each section with the answer, not the background.

Choose depth over breadth

A library of 120 thin pages routinely loses to a focused set of 12 pages carrying real depth, named sources, and evidence.

	Before	After
Volume	120 pages, about 200 words each	15 pages, 800 or more words each
Substance	Generic descriptions, no attribution	Outcomes data, named program directors
Maintenance	Published once, never updated	Updated quarterly with new results
Result	Generic, no depth signal, not cited	Cited across AI answers

Add images and video

AI surfaces multimedia from the pages it cites. Consider a newsroom that publishes a strong, original investigation: solid reporting, named sources, hard data, but no images or charts. A competitor's shorter story with one labeled chart wins the visual AI panel, while the original is cited only in text. Add relevant media to priority pages, give every image a descriptive label rather than a filename like "image123.jpg," and embed video where it helps.

Run a content gap analysis

Two methods reveal where you are losing.

- **Ask the AI.** Run your top queries in ChatGPT, Perplexity, and Google AI Mode and mark where you are not cited. Ask the AI what else people want to know to surface questions you never thought to answer, and note what each cited source has that you do not.
- **Study the competitor.** Identify who gets cited repeatedly, list the questions their pages answer that yours do not, and mark where you are thin, missing, or unstructured.

> KEY INSIGHT
>
> Close every gap to the same citable standard: answer first, labeled sections, named author, real evidence. Start with the gaps only your organization can fill from first-hand knowledge. Those carry the highest citation potential.

CHAPTER 06

Fix Area Three: Earn the Visit

Being cited is not the finish line. Once a citation sends a visitor to your page, the experience either reinforces your authority or erodes it.

Page experience still counts

Imagine a page cited in AI Overviews that takes 7.2 seconds to load on mobile, with most referred visitors leaving within eight seconds. Google watches post-click behavior, and citation authority erodes when the experience disappoints. Core Web Vitals, mobile usability, and HTTPS remain genuine ranking and retrieval factors. Run PageSpeed Insights on your most-cited pages, fix layout-shift and load-time issues, optimize images, and confirm HTTPS everywhere.

Right-size your schema

This is also where money quietly leaks. Consider three months spent building custom markup across every page specifically to improve AI visibility, with a citation rate after launch that is completely unchanged. There is no special markup that improves AI visibility. Structured data is worth keeping only where it earns standard Search rich results, and the budget aimed at "AI markup" is better redirected to content quality.

Lever	Do	Avoid
Performance	Fix LCP and CLS on cited pages; enable lazy loading	Ignoring mobile load time on high-traffic pages
Security	Confirm HTTPS sitewide	Mixed-content warnings that block citation credit
Schema	Keep markup that earns rich results in standard Search	Investing in "AI-specific" markup that does nothing

QUICK WIN

Audit your five most-cited pages for load speed this week. A cited page that loads slowly is a citation you are actively losing.

CHAPTER 07

What Google Always Rewards

Two fundamentals sit beneath every AI-search guideline and have rewarded organizations for years. They deserve deliberate attention.

Product and service knowledge

A service page with only a headline, a short paragraph, and a contact form leaves every buyer question unanswered (who it is for, what problem it solves, what the outcome looks like), so a competitor who answers those questions gets cited instead. Audit your highest-value service and product pages against four labeled questions: Who This Is For, What You Get, How It Works, What to Expect. Fill them with real outcomes and specific detail no competitor could copy, and refresh them quarterly.

Entity relationships

If an organization has operated in its sector for twenty years but a search by its name returns inconsistent descriptions across its own site, directory listings, and Google profile, Google cannot confidently identify what it is, and entity confidence is a prerequisite for citation. Make your name, description, and category consistent everywhere; add Organization schema to your homepage and About page; link internally from authority pages to the content you want cited; and earn mentions from credible sources in your sector.

Buyer Question	What a Citable Page Includes
Who is this for?	A specific audience, named plainly
What problem does it solve?	The concrete problem, in the reader's words
How does it work?	A clear, stepwise explanation
What should I expect?	Real outcomes, examples, and specifics

> KEY INSIGHT
>
> Entity confidence is built through consistency and association. Google trusts what the wider web describes the same way you describe yourself.

CHAPTER 08

Stop Funding the Noise

Every hour and dollar spent on tactics that do not work is one not spent on the fixes that do. Google was unusually direct about what to abandon.

Stop Doing This	Because
AI-specific schema markup	No structured data is required for generative AI
Rewriting content for an "AI tone"	Google rewards content written for people
AEO and GEO as separate workstreams	Optimizing for AI search is still SEO
More pages to raise AI visibility	Fewer deep pages outperform many thin ones
Chasing "guaranteed" AI Overview placement	Google does not guarantee inclusion, regardless of optimization
Building llms.txt files to feed AI	Crawlers treat them as ordinary text files; there is no special pathway
Pre-chunking content for AI	Google extracts the relevant passage from a well-structured page on its own

A single publishing test

Before any page goes live, ask four questions. Is this information only your organization can provide? Does it answer a specific question your audience would actually search for? Is there evidence: data, named people, real outcomes? Would a reader leave with something they could not find elsewhere? If the answer is no to any of them, the page is not ready.

The FAQ change

On May 7, 2026, Google removed FAQ rich results from search results entirely. The schema is finished as a visibility tactic, related reporting is being retired, and historical data should be exported before API support ends. But the content still matters: convert FAQ schema pages into structured content with labeled sections that answer real questions.

A note on agentic search

Emerging AI agents navigate pages through screenshots, the page's underlying structure, and the accessibility tree. Google's guidance is to explore this only if it is relevant and you have spare capacity. The reassuring part is that every fix in this ebook (structure, clarity, accessibility) builds agentic readiness automatically. It is the same investment, not a new one.

> **THE TEST THAT MATTERS**
>
> "Does this content provide substantial value compared to other pages in search results?" This is the question at the heart of Google's May 2026 guidance, and the only test that ultimately matters.

CHAPTER 09

Measure What Changed

Improvement you cannot see is improvement you cannot defend at budget time. Track weekly, screenshot everything, and build a baseline you can point to.

Fix Area	What to Measure	Tool
Indexability	Coverage report: excluded-pages count	Search Console
Crawl budget	Crawl Stats: pages crawled per day	Search Console
JS rendering	Crawled versus indexed gap	URL Inspection Tool
Canonical URLs	Indexed-URL count per key page	Search Console, Screaming Frog
Content extraction	Click-through rate on target pages	Search Console
Images and video	Impressions on priority pages	Search Console
Page experience	Core Web Vitals: LCP and CLS	PageSpeed Insights
Schema	Rich-result appearances	Rich Results Test
AI citation	Five priority queries, weekly	ChatGPT, Perplexity, AI Mode

The citation check

Once a week, search your five highest-value queries across the major AI tools, note which sources are cited, and screenshot the result. Where you are absent, identify what the cited source has that yours lacks, and route that finding back to the relevant content fix.

	Before	After (90 days, one page restructured)
Citation	Two competitors cited for the target query	Your site cited in three of five AI tools
The page	300 words, body field only	900 words, labeled sections, outcomes by cohort

KEY INSIGHT

A weekly citation check is the most honest measure of AI-search progress. Rankings hint at potential. Citations prove it.

CHAPTER 10

Your 90-Day Roadmap

Sequence matters. You cannot be cited if you cannot be found, so retrieval is fixed before any content investment.

Phase	Weeks	Focus
Diagnose	1 to 2	Search Console audit, crawl report, citation check on 10 priority queries
Fix retrieval	3 to 5	Indexability, crawl budget, canonical URLs, HTML delivery confirmed
Fix content	6 to 9	Five priority pages restructured: labeled sections, depth, first-hand authority
Measure and embed	10 to 12	Weekly citation check, AI Overview tracking, Core Web Vitals review

Five takeaways to carry forward

1. AI search is your audience's shortlist. Not cited means not considered, before anyone visits.
2. The gap is structural. You need content AI can retrieve and extract, not simply more content.
3. Retrievability comes first. Fix indexing before any content investment.
4. Depth, structure, and named authority are what AI cites. Thin, well-designed pages stay invisible.
5. Google confirmed what not to do: AI markup, FAQ schema, and volume publishing are wasted budget.

The mistakes that quietly cost citations

The Mistake	Why It Costs You
Generic content, no unique insight	Passes "publishable," fails "citable"
Thin product or service pages	No depth signal; a more specific competitor is cited
Body fields only, no structure	AI cannot extract an answer; the page is skipped
Duplicate content across URLs	Authority splits; no page earns enough signal
No author or organization credibility	Anonymous pages are deprioritized
Slow mobile load	Cited once, then dropped as visitors leave
No entity clarity	Google cannot confidently identify your organization

THE BOTTOM LINE

Your Drupal platform is natively capable of every fix in this ebook. You are not rebuilding. You are configuring.

YOUR NEXT STEP

Find Out Where You Stand

The shift from displaced to cited is not a trend to wait out. It is the new shape of discovery, and it rewards the organizations that act on the fundamentals first. The work is finite, sequenced, and, on Drupal, entirely achievable without a rebuild.

The fastest way to know where to begin is a focused audit of where your site stands today. A BinaryWorks AI Search Readiness Audit reviews:

- Google visibility across your key pages
- Content quality and depth assessment
- Page structure and extractability review
- Duplicate and canonical URL issues
- Mobile performance and Core Web Vitals
- Technical and structured-data markup review

That single snapshot tells you which of the fixes in this ebook will move your citations fastest.

Get your free AI Search Readiness Audit

A clear, prioritized view of what to fix first on your Drupal site.

Request Your Free Audit

ABOUT

About BinaryWorks

Your Digital and AI Growth Partner, trusted since 2009.

BinaryWorks partners with enterprise organizations to transform disconnected digital platforms into intelligent ecosystems that drive business growth. Combining CMS engineering on Drupal, WordPress, and Adobe Experience Manager with AI, data infrastructure, and performance marketing, we deliver end-to-end digital transformation across high-stakes industries. Built for complexity. Engineered for outcomes.

What we do

Capability	Focus
CMS and Digital Experience	Drupal, WordPress, and AEM development, migration, and support
AI, Automation, and Cloud	Strategy through production: LLMs, agents, automation, and cloud
AI Visibility and Search	Generative Engine Optimization (GEO) and Answer Engine Optimization (AEO)
Growth and Performance Marketing	Full-funnel programs measured by business impact
Experience Design	Research-driven, accessible interfaces
Security and Compliance	HIPAA, SOC 2, PCI-DSS, ADA, and WCAG built in from day one
Migration and Modernization	Zero-downtime migrations and legacy modernization

Industries we serve

Higher Education, Healthcare and Life Sciences, Nonprofits and Associations, Government and Public Sector, Finance, Manufacturing, Retail, Insurance, and Enterprise B2B.

Get in touch

Address	7000 Central Parkway, Suite 220, Atlanta, GA 30328
Phone	470-604-9885
Email	hello@thebinaryworks.com
Web	www.thebinaryworks.com

© 2026 BinaryWorks. All rights reserved.